

QUIC Observability

Rationale and Tools


Jana Iyengar

Fastly


jri@fastly.com

(with many thanks to Kazuho Oku)


What are we talking about?



The QUIC Standard



The QUIC Standard



Middleboxes

“[...] intermediary device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and destination host” - RFC 3234

Home routers (NATs)

Firewalls

Application load balancers (HTTP)

Protocol accelerators (PEPs)

Sometime, a while ago

First byte of gQUIC packet was *flags*

Sometime, a while ago

First byte of gQUIC packet was *flags*
: unencrypted, and had been 0x07 for a while

Sometime, a while ago

First byte of gQUIC packet was *flags*
: unencrypted, and had been 0x07 for a while

We flipped a bit.

Sometime, a while ago

First byte of gQUIC packet was *flags*
: unencrypted, and had been 0x07 for a while

We flipped a bit.

“users cannot reach any Google property over Chrome!”

What had happened

Firewall

- : allowed first packet in both directions
- : blackholed all subsequent packets

What had happened

Firewall

- : allowed first packet in both directions
- : blackholed all subsequent packets

“in wireshark, noticed that first byte was always the same”

What had happened

```
if udp_payload[0] == 7: QUIC
```

Protocol design maxim

"the ultimate defense of the end to end mode is end to end encryption"

David Clark, J. Wroclawski, K. Sollins, and R. Braden, *Tussle in Cyberspace: Defining Tomorrow's Internet*. IEEE/ACM ToN, 2005.

Current Status

Work at IETF for past 2 years

Strong focus on security and privacy

Network operator woes

Strong focus on avoiding ossification

Encryption

GREASEing

Several implementation efforts

Apple (ATS), Fastly (quicly/H2O), Facebook, Firefox, F5,
Google (Chromium), Microsoft, LiteSpeed, quic-go (Caddy)

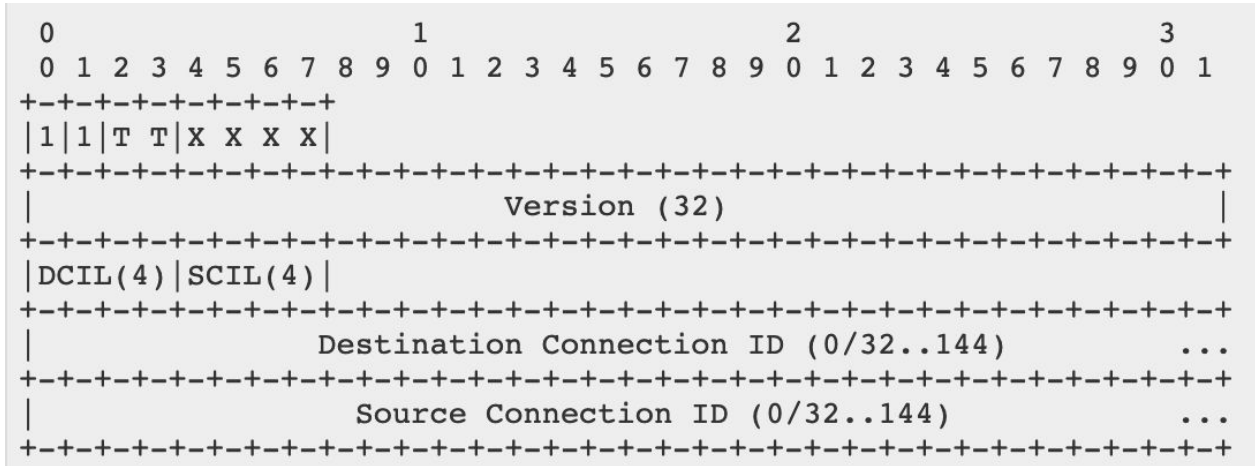
QUIC Packet Format

Long header

Short header

QUIC Packet Format

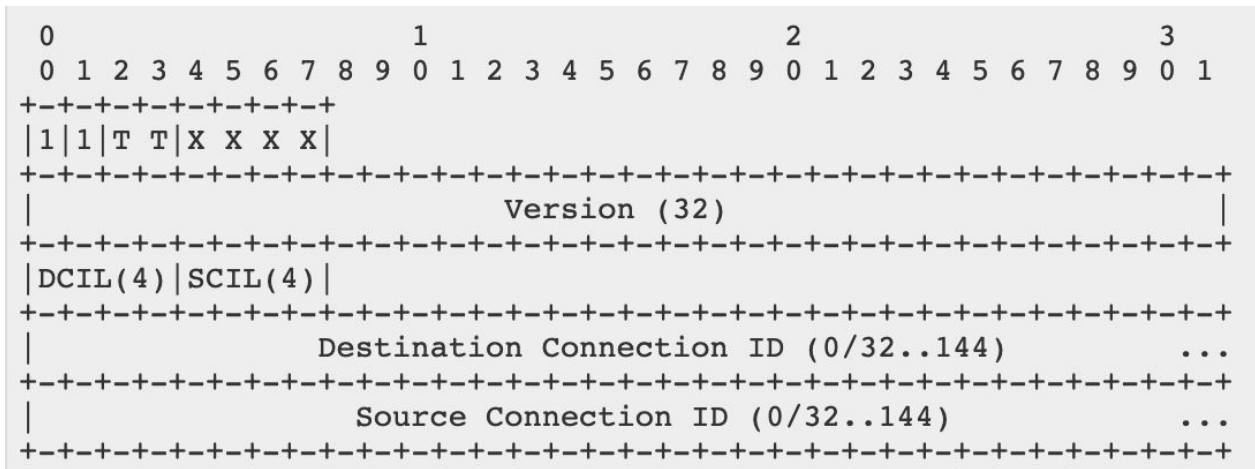
Long header



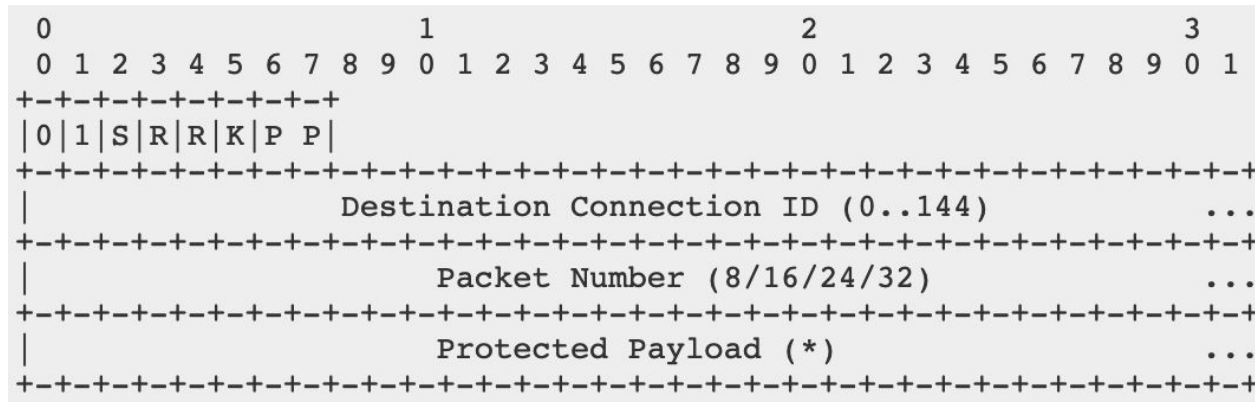
Short header

QUIC Packet Format


Long header





Short header



Frames





Frames



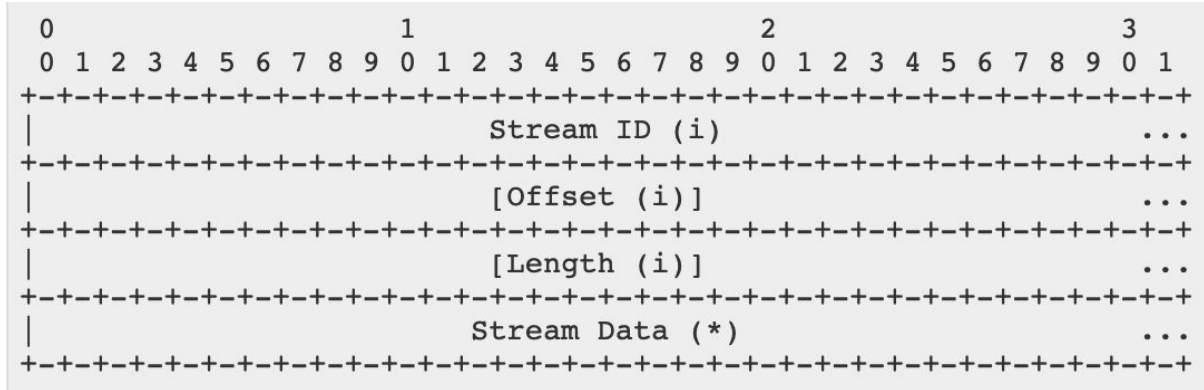
Type Value	Frame Type Name
0x00	PADDING
0x01	PING
0x02 - 0x03	ACK
0x04	RESET_STREAM
0x05	STOP_SENDING
0x06	CRYPTO
0x07	NEW_TOKEN
0x08 - 0x0f	STREAM
0x10	MAX_DATA
0x11	MAX_STREAM_DATA
0x12 - 0x13	MAX_STREAMS
0x14	DATA_BLOCKED
0x15	STREAM_DATA_BLOCKED
0x16 - 0x17	STREAMS_BLOCKED
0x18	NEW_CONNECTION_ID
0x19	RETIRE_CONNECTION_ID
0x1a	PATH_CHALLENGE
0x1b	PATH_RESPONSE
0x1c - 0x1d	CONNECTION_CLOSE

Frames

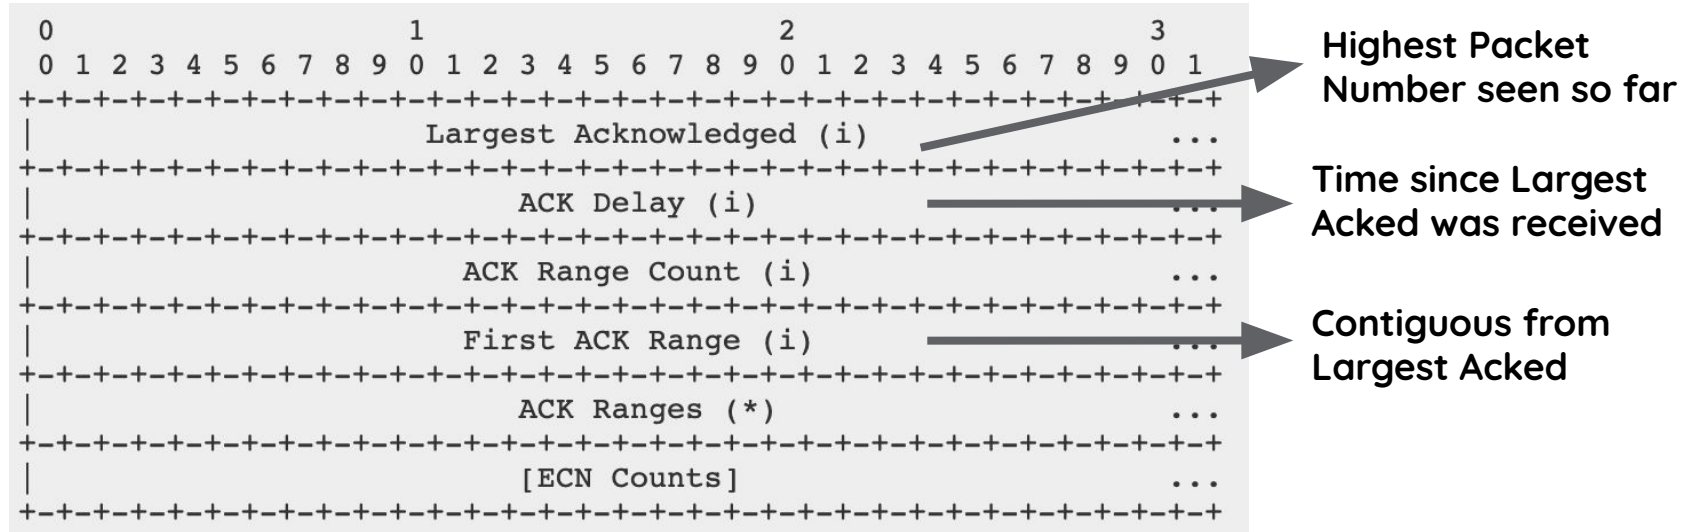


Type Value	Frame Type Name
0x00	PADDING
0x01	PING
0x02 - 0x03	ACK
0x04	RESET_STREAM
0x05	STOP_SENDING
0x06	CRYPTO
0x07	NEW_TOKEN
0x08 - 0x0f	STREAM
0x10	MAX_DATA
0x11	MAX_STREAM_DATA
0x12 - 0x13	MAX_STREAMS
0x14	DATA_BLOCKED
0x15	STREAM_DATA_BLOCKED
0x16 - 0x17	STREAMS_BLOCKED
0x18	NEW_CONNECTION_ID
0x19	RETIRE_CONNECTION_ID
0x1a	PATH_CHALLENGE
0x1b	PATH_RESPONSE
0x1c - 0x1d	CONNECTION_CLOSE

STREAM Frame



ACK Frame



QUIC Packetization: Example

QUIC Packet

Header = 0b01

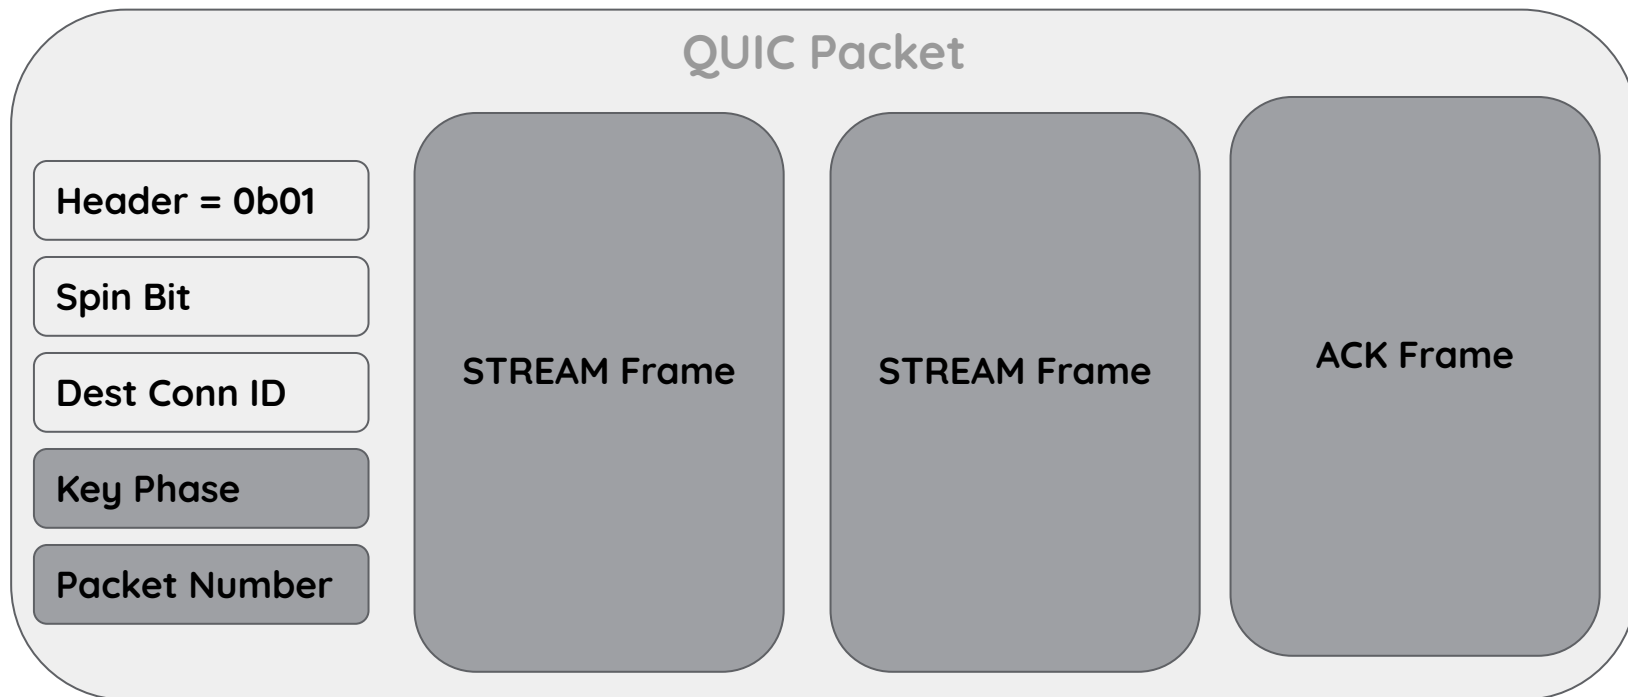
Spin Bit

Dest Conn ID

Key Phase

Packet Number

QUIC Packetization: Example



Tooling

In-network packet tracing

Wireshark dissector available

This isn't enough. Why?

Tooling

In-network packet tracing

Wireshark dissector available

This isn't enough. Why?

Endpoint-based packet tracing

Log packet and frame details at endpoint

(also log other transport info, such as cwnd)

Tooling

In-network packet tracing

Wireshark dissector available

This isn't enough. Why?

Endpoint-based packet tracing

Log packet and frame details at endpoint
(also log other transport info, such as cwnd)

quic-trace

QUICvis

Tooling: quic-trace

Written by Victor Vasiliev et al (Google)

Available at <https://github.com/google/quic-trace>

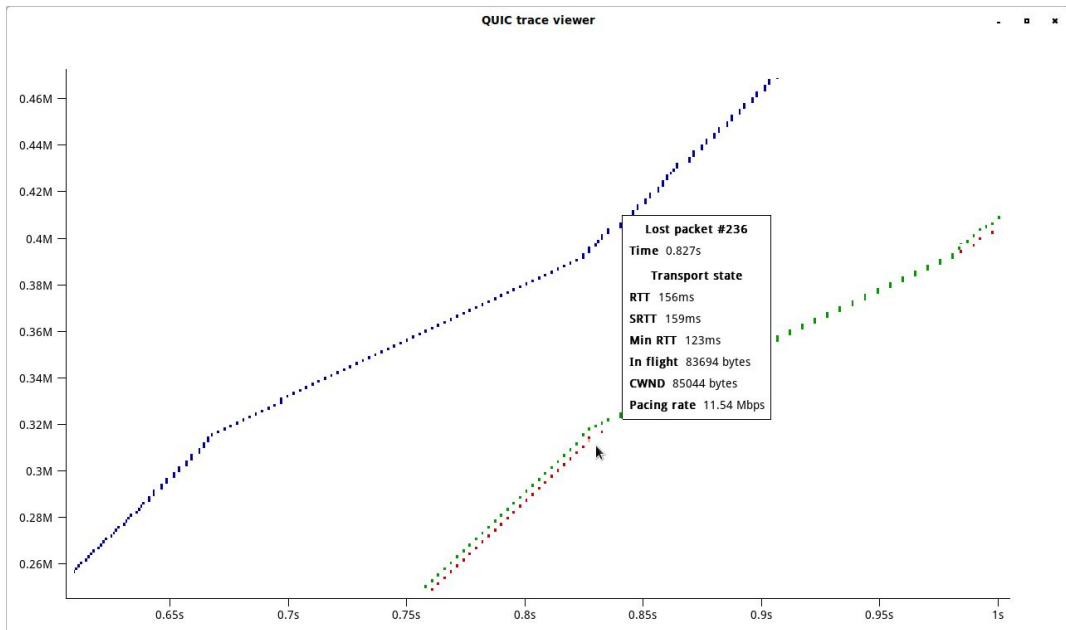
Input: protobuf or JSON

Tooling: quic-trace

Written by Victor Vasiliev et al (Google)

Available at <https://github.com/google/quic-trace>

Input: protobuf or JSON



Tooling: QUICvis

Written by Robin Marx et al

Available at <https://quic.edm.uhasselt.be/>

Input: JSON

Tooling: QUICvis

Written by Robin Marx et al

Available at <https://quic.edm.uhasselt.be/>

Input: JSON


```
1 {"connectionid": "0x763f8eaf61aa3ffe84270c0644bdbc2b0d", "starttime": 1543917600,
2  "fields":
3  ["time", "category", "type", "trigger", "data"],
4  "events": [
5  [50, "TLS", "0RTT_KEY", "PACKET_RX", {"key": ...}],
6  [51, "HTTP", "STREAM_OPEN", "PUSH", {"id": 0, "headers": ...}],
7  ...
8  [200, "TRANSPORT", "PACKET_RX", "STREAM", {"nr": 50, "contents": "GET /ping.html", ...}],
9  [201, "HTTP", "STREAM_OPEN", "GET", {"id": 16, "headers": ...}],
10 [201, "TRANSPORT", "STREAMFRAME_NEW", "PACKET_RX", {"id": 16, "contents": "pong", ...}],
11 [201, "TRANSPORT", "PACKET_NEW", "PACKET_RX", {"nr": 67, "frames": [16, ...], ...}],
12 [203, "RECOVERY", "PACKET_QUEUED", "CWND_EXCEEDED", {"nr": 67, "cwnd": 14600, ...}],
13 [250, "TRANSPORT", "ACK_NEW", "PACKET_RX", {"nr": 51, "acked": 60, ...}],
14 [251, "RECOVERY", "CWND_UPDATE", "ACK_NEW", {"nr": 51, "cwnd": 20780, ...}],
15 [252, "TRANSPORT", "PACKET_TX", "CWND_UPDATE", {"nr": 67, "frames": [16, ...], ...}],
16 ...
17 [1001, "RECOVERY", "LOSS_DETECTED", "ACK_NEW", {"nr": a, "frames": ...}],
18 [2002, "RECOVERY", "PACKET_NEW", "EARLY_RETRANS", {"nr": x, "frames": ...}],
19 [3003, "RECOVERY", "PACKET_NEW", "TAIL_LOSS_PROBE", {"nr": y, "frames": ...}],
20 [4004, "RECOVERY", "PACKET_NEW", "TIMEOUT", {"nr": z, "frames": ...}]
21 ]}
```


Tooling: QUICvis

Written by Robin Marx et al

Available at <https://quic.edm.uhasselt.be/>

Input: JSON






Tooling: QUICvis

Written by Robin Marx et al

Available at <https://quic.edm.uhasselt.be/>

Input: JSON

```
1 {"connectionid": "0x763f8eaf61aa3ffe84270c064"}
2 "fields":
3 ["time", "category", "type",
4 "events": [
5 [50, 0]
6 [51, 0]
7 ...
8 [20, 0]
9 [20, 0]
10 [20, 0]
11 [20, 0]
12 [20, 0]
13 [25, 0]
14 [25, 0]
15 [25, 0]
16 ...
17 [10, 0]
18 [20, 0]
19 [30, 0]
20 [40, 0]
21 ]}]
```



Tooling: QUICvis

Written by Robin Marx et al

Available at <https://quic.edm.uhasselt.be/>

Input: JSON

```
1 {"connectionid": "0x763f8eaf61aa3ffe84270c064"}
2 "fields":
3   ["time", "category", "type",
4   "events": [
5     [50, 0],
6     [51, 0],
7     [20, 0],
8     [20, 0],
9     [20, 0],
10    [20, 0],
11    [20, 0],
12    [20, 0],
13    [25, 0],
14    [25, 0],
15    [25, 0],
16    [10, 0],
17    [20, 0],
18    [30, 0],
19    [40, 0],
20    [40, 0],
21  ]
}
```

